



AC

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ :	A1	(11) International Publication Number:	WO 99/12381
H04Q 11/04, H04L 12/56		(43) International Publication Date:	11 March 1999 (11.03.99)

(21) International Application Number: PCT/US98/18582

(22) International Filing Date: 4 September 1998 (04.09.98)

(30) Priority Data:
08/924,817 5 September 1997 (05.09.97) US

(71) Applicant (for all designated States except US): OBJECTIVE COMMUNICATIONS, INC. [US/US]; 75 Rochester Avenue, Portsmouth, NH 03801 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): ROGERS, Steven [US/US]; Objective Communications, Inc., 75 Rochester Avenue, Portsmouth, NH 03801 (US).

(74) Agents: BANNER, Pamela, I. et al.; Banner & Witcoff, Ltd., 11th floor, 1001 G Street, N.W., Washington, DC 20001-4597 (US).

(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published

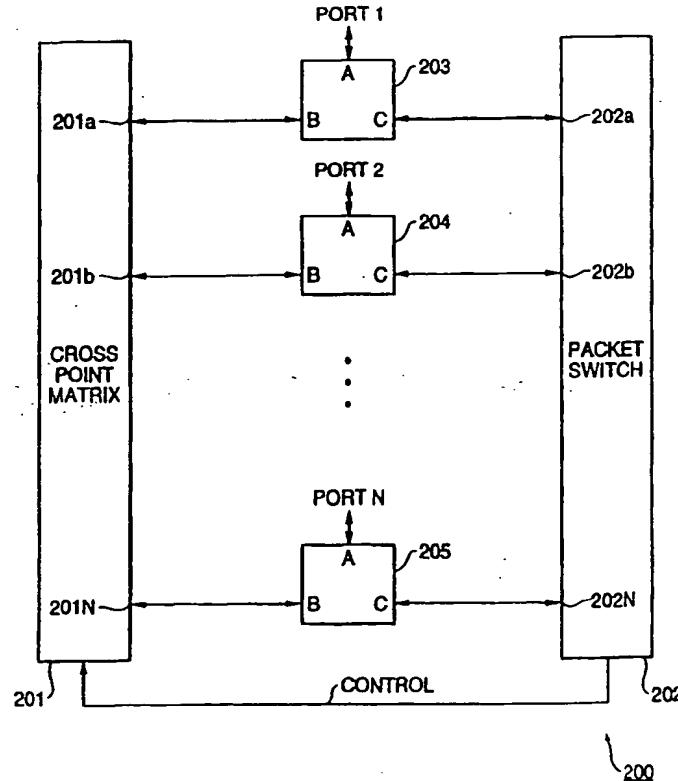
With international search report.

Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.

(54) Title: SCALABLE HIGH SPEED PACKET SWITCH USING PACKET DIVERSION THROUGH DEDICATED CHANNELS

(57) Abstract

A scalable high-speed packet switching device includes a packet switch with ports, a crosspoint matrix, and a plurality of subswitches, one for each port of the packet switch. A processor in the packet switch detects a high level of traffic between two ports and causes the subswitches and crosspoint matrix to divert packets away from the packet switch and through a dedicated channel in the crosspoint matrix.



BEST AVAILABLE COPY

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KR	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		

**SCALABLE HIGH SPEED PACKET SWITCH
USING PACKET DIVERSION THROUGH DEDICATED CHANNELS**

Inventor: Steven Rogers

5 BACKGROUND OF THE INVENTION

1. Technical Field

This invention relates generally to packet switching devices, and more particularly to a packet switch which can be scaled up to handle a large number of ports without a corresponding decrease in performance. The invention provides a technique for detecting high traffic levels between two ports and diverting packets between such ports through a dedicated channel.

10 **2.** Related Information

Packet switching techniques are conventionally used to transmit various types of digital data such as digitized voice, computer data, and video signals. Data (e.g., a telephone voice signal) is broken up into packets each of a specified size, wherein each packet typically includes, in addition to the digital data, a header indicating a destination address to which the packet should be sent. A packet switching network made up of packet switches can be used to route packets to their respective destinations without delay. See, for example, U.S. Patent No. 5,544,160 to Cloonan et al., entitled "Terabit per Second Packet Switch".

20
25
Packet switches may also be used in smaller communication networks to route data between devices. As one example, an interoffice videoconferencing system may have a requirement to selectively transmit digitized video and audio information to one or more recipients through a centralized switch. A packet switch can be designed to handle such digitized information, so that packets comprising a video signal from one video conference participant are switched to all participants of the video conference.

As the size of packet switching networks has grown, the complexity and processing requirements of packet switching devices has increased. For example, the

- 2 -

number of ports required for a particular packet switch has increased, requiring faster processors to handle the increased traffic.

Conventional packet switches (i.e., devices which route packets of digital data between two sets of ports) are well known. FIG. 1 shows in simplified form a conventional packet switch architecture including a CPU 101, a routing table 102, buffer memory 103, data bus 104, and packet interfaces 105 through 111 each corresponding to a port. Packets arrive at a port (e.g., port 1) and are routed to a destination port (e.g., port 2) based on information contained in packet headers.

Most packet switches rely on fast software to route packets. As shown in FIG. 1, for example, incoming packets are received by packet interfaces 105 through 111 and interpreted by CPU 101. CPU 101 temporarily stores incoming packets into buffer memory 103 and examines each packet header to determine which port the packet should be routed to for output. This examination usually involves searching a look-up table of addresses versus ports stored in routing table 102. The table includes a list of addresses and ports which are most advantageous to route the packet for a given address. Following the look-up operation, the packet is placed in the correct packet interface and sent to the destination port. Special messages can be transmitted to the packet switch to change routing table 102, thus changing the network routing parameters.

One problem with the conventional architecture of FIG. 1 is that it can take a long time to sort through routing table 102. As a network grows, the number of potential distant packet addresses grows. In fact, it may be that the packet routing table can be billions of addresses long, severely bogging down CPU 101.

The time it takes to receive and store packets while the routing process is underway will limit the number of ports that a packet switch can support. It also limits the number of packets that can be routed per second. A given CPU has a finite

- 3 -

processor bandwidth Z. This bandwidth will allow the CPU to successfully route X packets per second. Each packet interface has a bandwidth or maximum packet speed Y in packets per second. Thus, the number of ports N that can be successfully serviced by this type of packet switch is limited to $N(\max) = X/Y$.

5 A packet switch based on the architecture of FIG. 1 can only be scaled as far as the CPU bandwidth is scaled. This is difficult, as there is a limit on the processing bandwidth of a single CPU. Consequently, packet switches based on the general architecture shown in FIG. 1 cannot be easily scaled to support a large number of ports or addresses.

10 **SUMMARY OF THE INVENTION**

The present invention overcomes the aforementioned problems by providing a scalable packet switch which can be thought of as a hybrid between a "connectionless" switch and a "connectioned" switch. The architecture can scale to a much larger number of ports, and is not limited by the bandwidth of a single CPU. 15 The scalable packet switch, in response to detecting a high traffic rate between two node addresses (and hence between switch ports), creates a dedicated path between the ports for packets between those nodes, thus diverting the packets away from a bottleneck in the switch.

20 In one embodiment, the packet switch includes a conventional packet switch which is augmented with an additional subswitch per port and a crosspoint matrix. Each subswitch is designed to switch packets to one of only two possible destinations: the conventional packet switch or to the crosspoint matrix. By limiting each subswitch to only three ports, the processor bandwidth can be fully used, thus allowing the packet per second rate per port to be absolutely maximized.

25 The new scalable switch takes advantage of a well-known property of packet communications: the majority of the packets sent through a switch are sent as part of

- 4 -

a stream of packets. In other words, an application using packet transmission will usually engage in a session with another location served by a particular port. This session will continue for some time, with most of the traffic being transmitted between the two distant ends. When the session is over, the two parties will then move on to other sessions, between other distant parties. During each session, however, there is little packet traffic that is not transmitted between the two parties and thus between two particular ports.

In accordance with the invention, a dedicated connection is provided to transmit certain packets directly between ports (avoiding the main packet switch) while still allowing other packets to travel through the main packet switch.

Additional features and advantages of the present invention will become apparent through the following detailed description, the figures, and the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a conventional packet switch architecture which uses a CPU to look up address/port entries in a routing table.

FIG. 2 shows a packet switch in accordance with one embodiment of the present invention, including a crosspoint matrix 201, a plurality of packet subswitches 203 through 205 and a main packet switch 202.

FIG. 3 shows one embodiment of a packet subswitch of the type used in the system of FIG. 2.

FIG. 4 shows a video conferencing system employing a high speed packet switch to route video and audio packets among conference participants.

FIG. 5 shows a method for dynamically reconfiguring a crosspoint matrix to provide a dedicated path between high-traffic ports.

FIG. 6 shows a method for routing packets at each subswitch.

- 5 -

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

FIG. 2 shows a packet switch which employs various principles of the present invention. Packet switch 200 includes a conventional packet switch 202 which is augmented with a crosspoint matrix 201 and an additional subswitch (203 through 205) per port. In contrast to FIG. 1, where switch ports are arranged across the top of the figure, switch ports in FIG. 2 are arranged down the middle of the figure.

In the inventive architecture of FIG. 2, each sub-switch is designed to switch packets with only two possible destinations. A packet arriving at a port (e.g., PORT 1) can either be routed to the main packet switch 202 or to crosspoint matrix 201. In accordance with the invention, packets flowing between high-traffic ports are diverted through crosspoint matrix 201 using a dedicated path instead of going through packet switch 200. Each sub-switch 203 through 205 may comprise a small standard packet switch with $N=3$ ports. The operation of FIG. 2 is explained in more detail below.

FIG. 3 shows one possible embodiment for subswitch 203 of FIG. 2. By limiting each subswitch to exactly three ports, the processor bandwidth can be fully utilized, thus allowing the packet per second rate per port to be absolutely maximized. As shown in FIG. 3, each subswitch includes a processor such as a CPU 301, a subswitch routing table 302, a buffer memory 303, data bus 304, and three port interfaces 305 through 307. A first port interface 305 provides the external packet switch port connection (i.e., this replaces port interface 105 of FIG. 1). A second port interface 306 couples the CPU to the cross point matrix. A third port interface 307 couples the CPU to the main packet switch (i.e., packet switch 202 of FIG. 2).

Subswitch routing table 302 and buffer memory 303 may reside in different areas within a single computer memory. CPU 301 may comprise any of various commercially available processors, such as 8-bit or 16-bit processors. Alternatively, CPU 301 may comprise a portion of an application-specific integrated circuit, or may

- 6 -

be constructed of other well known logic functional units. The essential function of CPU 301 is to route packets between PORT 1 and the crosspoint matrix or the main packet switch based on a configuration command.

5 Normally, subswitch 203 merely transfers incoming packets from PORT 1 (through interface 305) to main packet switch 202 (through interface 307) and vice versa. However, subswitch 203 can be commanded to create an entry in subswitch routing table 302 such that all packets destined for a particular address are instead diverted to crosspoint matrix 201 (via port interface 306) and vice versa. Commands to reconfigure subswitch 203 may be provided either through crosspoint matrix 201 10 or from packet switch 202 via port interface 307. The subswitch may be commanded to divert packets destined for a particular logical address, physical address, node address, port, or any other designation.

15 The inventive packet switch relies on the fact that when an application such as video conferencing transmits information between 2 locations, a session is created, causing a large number of packets to flow during the session. This session will continue for some time, with most of the traffic being transmitted between the two distant ends. When the session is over, the two parties will then move on to other sessions, between other distant parties. During the session, there is very little packet traffic that is not transmitted between the two parties.

20 One example of such a session is a multimedia communications session. If a user is sending a series of video or graphics images, then the data will exhibit a high packet volume directed at the distant end with little or no packets sent anywhere else. During this session, a "virtual" connection is established between the two parties.

25 FIG. 4 shows a videoconferencing system using a high speed packet switch according to various aspects of the invention. High speed packet switch 450 includes six ports each of which is coupled through appropriate digital conversion hardware

- 7 -

5

(e.g., A/Ds and D/As 457 through 462) to video conferencing stations 451, 452 and 453. Such connections may be made over a LAN, telephone wires, or other suitable media. High speed packet switch 450 also includes a traffic rate comparator 450a which may be implemented in hardware, software, or a combination of the two. In one embodiment, traffic rate comparator compares traffic levels between addresses flowing through the switch, determines whether a threshold has been exceeded, and causes packets to be diverted through a crosspoint matrix when traffic levels are exceeded between two ports corresponding to the high-traffic addresses.

10

Each video conferencing station may comprise a personal computer with a color display, audio equipment including a microphone and speakers, and a video camera such as camera 454. It will be appreciated that other components may also be included in such a video conferencing system to provide various features such as camera control, filtering, conference bridging/mixing, etc.

15

20

Video and audio information from each conference station is digitized and transmitted to high speed packet switch 450, which is constructed in accordance with the architecture of FIGs. 2 and 3. Packet switch 450 examines each incoming packet and routes the packet to its destination port. For example, if conference station 453 is conducting a video teleconference with conference station 451, packet switch 450 routes data packets containing video and audio information from station 453 (received on port 1) to station 451 (through port 6). Similarly, packet switch 450 routes data packets containing video and audio information from station 451 (received on port 5) to station 453 (through port 2). It will be appreciated that bi-directional ports may be used instead of unidirectional ports.

25

In contrast to conventional systems, however, packet switch 450 detects the condition that a large traffic volume is flowing between the respective ports, and reconfigures subswitches within the packet switch to transmit data through the

- 8 -

crosspoint matrix rather than looking up each address in the main packet switch. Accordingly, increased traffic volume can be handled.

Referring again to FIG. 2, processing requirements on the main packet switch 202 are greatly reduced by diverting packets destined for a particular port through crosspoint matrix 201. Crosspoint matrix 201, which may comprise a digital crosspoint matrix, can establish a direct link between two of the sub-switch ports in the overall switch. This matrix can be controlled by the CPU in main packet switch 202 over a control line CONTROL. When main packet switch 202 detects a high traffic rate between two ports, it controls crosspoint matrix 201 to establish a direct connection between the appropriate subswitches. Main switch 202 informs, through special control packets or via direct control signals, the subswitches to route the connected traffic for those addresses directly through the crosspoint direct connection instead of through main switch 202. In the case where data is not part of the "virtual circuit", it is directed to the main packet switch and routed in the normal manner.

Crosspoint matrix 201 may comprise any of various types of crosspoint switches which provide a dedicated path between two designated ports. Such crosspoint switches are well known and no further elaboration is needed. Packet switch 202 may comprise any of various types of packet switches which route incoming packets to one or more destination ports based on address information in a packet header (e.g., see FIG. 1). In contrast to conventional packet switches, however, packet switch 202 may include a CPU which detects a condition that a large number of packets are flowing between two addresses (or ports) and generates a signal on control line CONTROL to configure crosspoint matrix 201 to provide dedicated paths between two corresponding subswitches. Additionally, the corresponding subswitches are commanded to divert incoming and outgoing packets for the particular addresses through crosspoint matrix 201 rather than through packet switch 200. Each subswitch

can accomplish this function by creating an entry in its corresponding subswitch routing table to reflect new port assignments.

Diverting traffic through crosspoint matrix 201 offloads main packet switch 202 from making repetitive packet routing. In some systems, such as multi-media communications systems, 90% or more of the packet traffic would be associated with these virtual connections. In this case the crosspoint matrix acts as a high-speed accelerator. In many systems of this kind the connection-oriented packet traffic is 99.99% or more of the total.

FIG. 5 shows a method for detecting traffic patterns and reconfiguring the packet switch according to one aspect of the present invention. The method may be implemented, for example, in a conventional packet switch by reprogramming the CPU to monitor traffic and issue commands accordingly. Alternatively, separate circuitry can be provided to detect traffic patterns and configure the crosspoint matrix and port subswitches. It may also be possible to implement the traffic monitoring function directly in the subswitches, such that the subswitches themselves detect traffic patterns and issue commands to reconfigure the crosspoint matrix and subswitch routing tables. All of these possibilities are well within the scope of the present invention.

Beginning with step 501, the traffic level between pairs of addresses or ports is monitored. This may be accomplished in any of various ways, such as creating an entry for each pair of source/destination addresses and incrementing a packet counter for each packet which passes through the pair corresponding to the entry. Once the address has been established, a port number can be established or determined.

In step 502, a test is made to determine whether the traffic level between the two ports exceeds a threshold. The threshold may be statically created or may be dynamically reconfigured depending on the circumstances (i.e., select the highest 10

- 10 -

traffic ports; select all ports which exceed a predetermined limit; or use statistical sampling techniques). If the threshold has not been exceeded, the monitoring continues in step 501.

5 If a traffic threshold has been exceeded, then in step 503 the crosspoint matrix is reconfigured to create a dedicated path between the two ports for the specified addresses. Additionally, in step 504, the subswitches associated with each port are commanded to divert traffic for the specified address. In one embodiment, when traffic levels fall below a specified minimum, the diversion of packets may be reversed.

10 Suppose that three nodes in a system are communicating (nodes 1, 2, and 3). Assume that nodes 1 and 2 are conducting a video conference or other multimedia session which involves a large quantity of traffic, and node 3 only periodically transmits a message to the other two nodes. In accordance with one aspect of the invention, a large traffic volume between nodes 1 and 2 would be detected, and a 15 dedicated path would be established for communicating between ports for those nodes by commanding two subswitches corresponding to the ports to divert packets going to or coming from the specified nodes through the crosspoint matrix. However, packets going to or coming from node 3 would not be diverted through the crosspoint matrix, and would instead travel through the packet switch according to conventional packet 20 switching techniques.

25 FIG. 6 shows a method for routing packets at each subswitch. Beginning in step 601, each subswitch processor looks up an address extracted from an incoming packet header in the local subswitch routing table. In step 602, if a tag associated with that address indicates that the packet should be routed to the crosspoint matrix, then in step 603 the packet is routed to the crosspoint matrix and hence directly to another subswitch. If, on the other hand, the tag or port assignment indicates that the packet

- 11 -

should not be routed to the crosspoint matrix, then in step 604 the subswitch routes the packet to the main packet switch where it is processed normally.

A switch constructed in accordance with the inventive principles is scalable since each sub-switch is associated with a port (i.e., N ports requires N sub-switches).

5 The crosspoint matrix also grows with the number of ports and can be implemented in digital hardware. Although it scales in complexity as N^2 , it is feasible to build large crosspoint matrix devices of 100 x 100 in a single integrated circuit.

There is no speed penalty in the sub-switch as the number of ports increases. For virtual circuit data the latency through an input subswitch, to the output port 10 subswitch, is the same for a small number of ports as well as for a large number of ports. In cases where the main packet switch is off-loaded by 90% the size of the over-all switch can now grow by a factor of ten because the main packet switch has one tenth the traffic it had before. If the connection oriented traffic is 99% of the total, then the scalable switch can grow to one hundred times the size of a standard 15 packet switch, etc.

The result of implementing a packet switch in accordance with the present invention is that with the same CPU type a switch can be implemented that has many more connection ports. It can scale, depending on the percentage of traffic that is connection-oriented, to ten or even one hundred or more times the number of ports on 20 a standard switch.

Another advantage of the scalable switch design is that it reduces the average latency through the switch. At first this would not be the case, since for a single isolated packet the route through the scalable switch would have to go through three switches, the main switch and two sub-switches. However, for connection oriented 25 packets the latency delay in the main switch can very large, depending on the search algorithm. With the scalable switch the latency is limited to the latency of two sub-

- 12 -

switches. In addition, the two subswitches have only to make a simple two-way routing decision. In fact, the subswitch, on the input side, need only route packets by comparing them to a single destination address. That address is the distant destination in the virtual connection. All other packets go to the main switch.

5 The one-address routing scheme lends itself to a common speed-up technique called "cut-through" addressing. With cut-through routing, the subswitch looks for the address field as the packet is being received and begins transmitting the packet to the proper port even before the whole packet is received. Because only one address is searched for, it is possible to build a very efficient cut-through switch. This
10 minimizes the total latency even further.

15 Finally, on the output sub-switch, the output packet from the main switch or from the cross-point matrix is always routed to the output port. These packets are assumed to be targeted for the distant connection. In this case, the outbound packet can be very quickly routed, using an almost zero latency cut-through technique. The only reason for buffering the outbound packet is to prevent a collision from the other outbound route. For example, if a packet is arriving to the output sub-switch from the crosspoint matrix at the same time as a packet from the main switch, then the second packet must be buffered while the first packet is forwarded.

20 The advantages of the scalable switch are many. First, the number of switch ports possible with a given CPU bandwidth is greatly increased. Conversely, the data rate for each port in the switch can be significantly increased. Finally, the latency in the switch, on average, can be significantly reduced for the connection-oriented traffic.

25 It is apparent that many modifications and variations of the present invention are possible, and references to specific values are by example only. Switch ports, for example, can be either unidirectional or bidirectional. Various functions can be

- 13 -

implemented in either hardware or software or a combination of both.

5

Reference numerals in the appended method claims identifying steps are for convenience only and are not intended to imply a necessary ordering of the steps. It is apparent that the method steps of the invention may be practiced in a different ordered sequence from that illustrated without departing from the scope of the invention. It is, therefore, to be understood that within the scope of the appended claims the invention may be practiced otherwise than as specifically described.

CLAIMS

1. An improved packet switching device, comprising:

5 a packet switch comprising a plurality of switch ports, wherein the packet switch receives packets at each switch port, determines a destination switch port for each packet by comparing an address in each packet with packet switch routing information, and routes each packet to the destination switch port;

10 a crosspoint matrix comprising a plurality of matrix ports, wherein the crosspoint matrix is configurable to provide direct paths between any two of the matrix ports; and

15 a plurality of subswitches, each coupled to the packet switch and to the crosspoint matrix, each subswitch including a first subswitch port at which packets are received or transmitted, a second subswitch port coupled to the crosspoint matrix, and a third subswitch port coupled to one of the plurality of switch ports of the packet switch;

20 15 wherein each subswitch is configurable to divert certain packets away from the packet switch to the crosspoint matrix and to allow other packets to pass through to the packet switch.

25 20 2. The improved packet switching device of claim 1, wherein each subswitch comprises a processor and a subswitch routing table, wherein the processor compares an address of each packet arriving at the subswitch with subswitch routing information stored in the subswitch routing table and routes packets having an address corresponding to the crosspoint matrix to the crosspoint matrix and routes packets having an address corresponding to the packet switch to the packet switch.

25 25 3. The improved packet switching device of claim 2, wherein the processor is responsive to a subswitch configuration command to change the subswitch routing information stored in the subswitch routing table.

- 15 -

4. The improved packet switching device of claim 3, wherein the processor receives the subswitch configuration command from the packet switch in response to a traffic rate comparison between two of the plurality of packet switch ports.

5. The improved packet switching device of claim 3, wherein the processor receives the subswitch configuration command from the crosspoint matrix.

6. The improved packet switching device of claim 1, wherein each of the switching ports of the packet switch are bidirectional.

10 7. The improved packet switching device of claim 1, wherein the packet switch performs a traffic rate comparison for addresses extracted from packet headers received at the packet switch, and, responsive to determining that a traffic rate between two addresses exceeds a threshold, transmitting a configuration command to two of the plurality of subswitches each corresponding to one of the packet switch ports which exhibits the traffic rate which exceeds the threshold to divert subsequent packets containing the two addresses through the crosspoint matrix.

15 8. The improved packet switching device of claim 7, wherein the packet switch further issues a command to the crosspoint matrix to set up a dedicated path between two of the plurality of subswitches.

20 9. The improved packet switching device of claim 1, wherein the packets contain packetized video and audio information corresponding to a video teleconference.

25 10. A system including the improved packet switching device of claim 1 and a plurality of video conference stations each of which transmits and receives packetized video information for a conference participant, wherein each conference station is coupled to at least one of the first subswitch ports, and wherein during a video conference between two of the video conference stations, the improved packet switching device detects a high traffic rate between the two video conference stations

- 16 -

and causes subsequent packets between the two video conference stations to be diverted through the crosspoint switch.

11. An improved packet switching device, comprising:

5 a packet switch comprising a plurality of switch ports each including a port interface for transmitting and receiving packets, a packet switch memory including routing information which associates packet addresses with switch ports, and a processor which extracts packet addresses from packets received at the switch ports, looks up each packet address in the packet switch memory, determines a destination switch port for each packet address on the basis of the looked up packet address, and transmits each packet received at the switch ports to a respective destination switch port;

10 a crosspoint matrix including a plurality of matrix ports, wherein the crosspoint matrix is configurable to provide direct paths between any two of the matrix ports; and

15 a plurality of subswitches, each subswitch comprising a first port at which packets are received or transmitted, a second subswitch port coupled to one of the plurality of matrix ports of the crosspoint matrix, a third subswitch port coupled to one of the plurality of switch ports of the packet switch, a subswitch routing memory which associates packet addresses with either the second subswitch port or the third subswitch port, and a processor which extracts addresses from packets received at the first subswitch port, compares the extracted addresses received from the first subswitch port with packet addressees in the subswitch routing memory, and transmits packets received from the first subswitch port to the second subswitch port for packets having extracted addresses associated with the second subswitch port in the subswitch routing memory and transmits packets received from the first subswitch port to the third subswitch port for packets having extracted addresses associated with the third subswitch port in the subswitch routing memory.

- 17 -

12. The improved switching device of claim 11, wherein the packet switch identifies high traffic levels between two of the packet switch ports and, in response thereto, causes the crosspoint matrix to provide a dedicated path between two of the plurality of subswitches and causes the two subswitches to divert subsequent packets away from the packet switch and through the dedicated path in the crosspoint matrix.

5 13. The improved switching device of claim 12, wherein the packet switch compares a traffic rate between two packet addresses extracted from packets flowing between two of the plurality of switch ports with a threshold.

10 14. The improved switching device of claim 11, further comprising a traffic comparator circuit which, in response to determining that a high traffic rate exists between two of the packet switch ports, configures the crosspoint matrix to provide a dedicated path between two of the plurality of subswitches.

15 15. The improved switching device of claim 14, wherein the traffic comparator circuit configures the two subswitches to divert packets having an address corresponding to the high traffic rate through the dedicated path of the crosspoint matrix.

16. A method of improving performance in a packet switching device having a plurality of ports through which packets are switched, the method comprising the steps of:

20 (1) detecting a condition that a high traffic rate exists between two of the plurality of ports;

25 (2) in response to detecting the high traffic rate, providing a dedicated path which diverts packets having addresses associated with the high traffic rate away from the plurality of ports, wherein such packets are routed to their destination through the dedicated path instead of through the plurality of ports; and

(3) continuing to route packets which do not have addresses associated with the

- 18 -

high traffic rate through the plurality of ports.

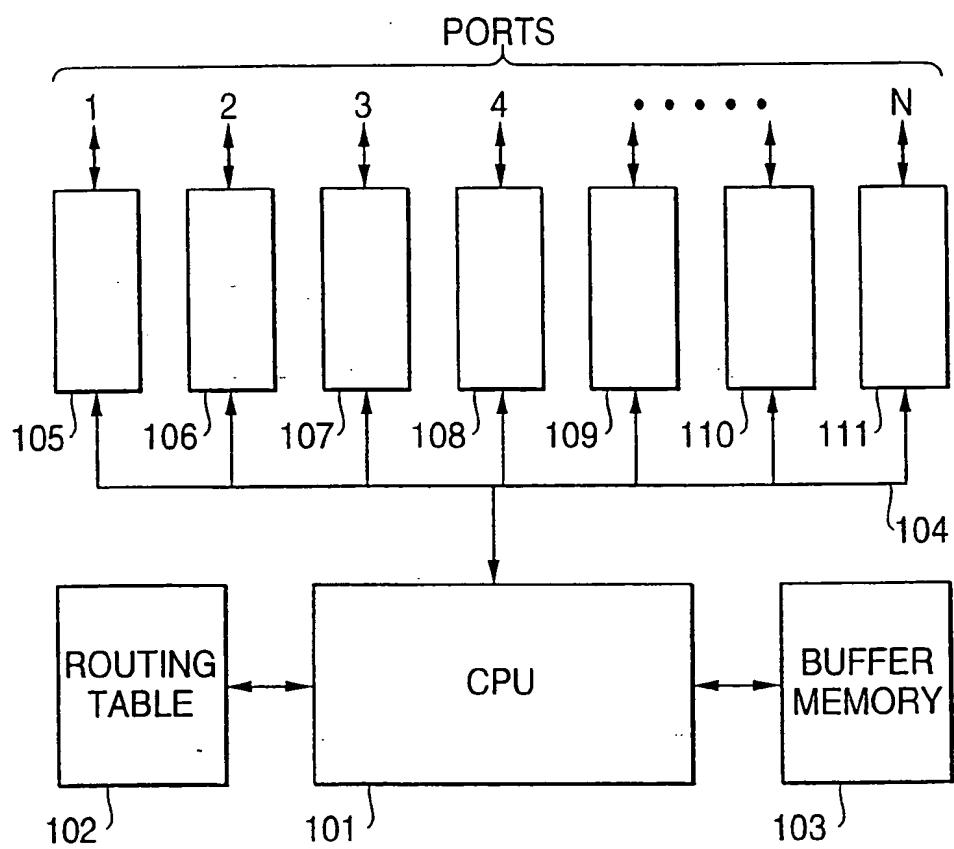
17. The method of claim 16, wherein step (2) comprises the step of configuring a crosspoint matrix, and two subswitches coupled between the crosspoint matrix and to two of the ports, to divert packets through the subswitches to the crosspoint matrix.

5 18. The method of claim 17, wherein step (2) comprises the step of using a processor and a subswitch routing table in each of the two subswitches to divert the packets.

10 19. The method of claim 18, wherein step (2) comprises the step of diverting packets containing digitized audio and video information associated with a video teleconference.

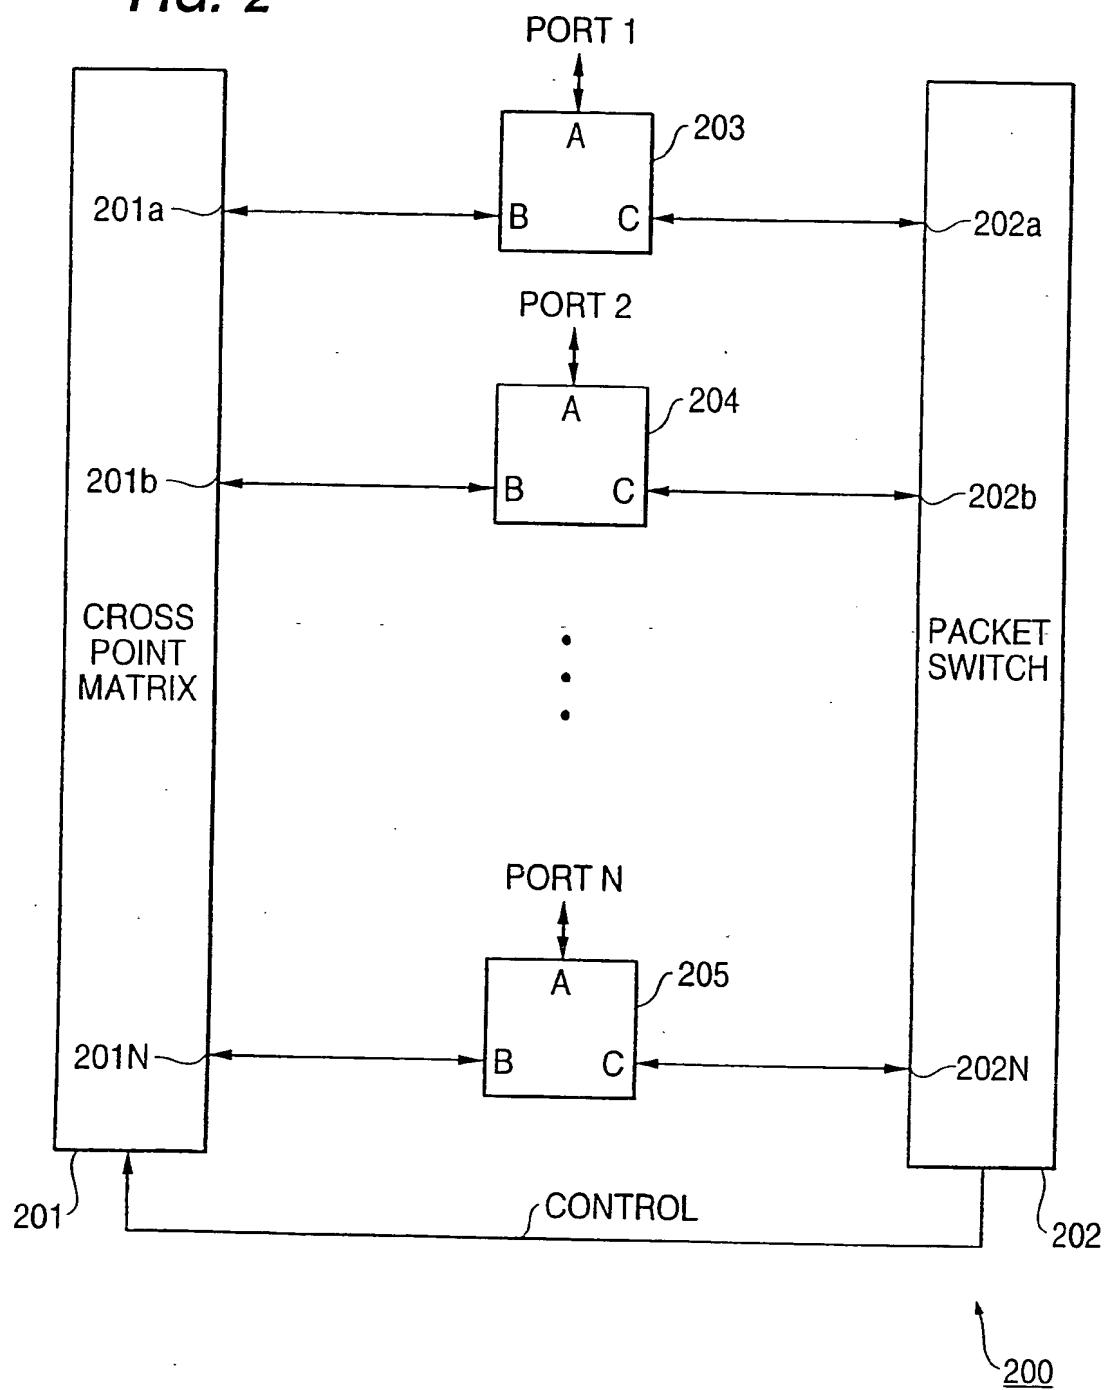
15 20. The method of claim 18, wherein step (1) comprises the step of using a traffic comparison circuit within the packet switch to detect the high traffic rate, and wherein step (2) comprises the step of issuing a command from the packet switch to a crosspoint matrix to provide the dedicated path.

FIG. 1
(PRIOR ART)



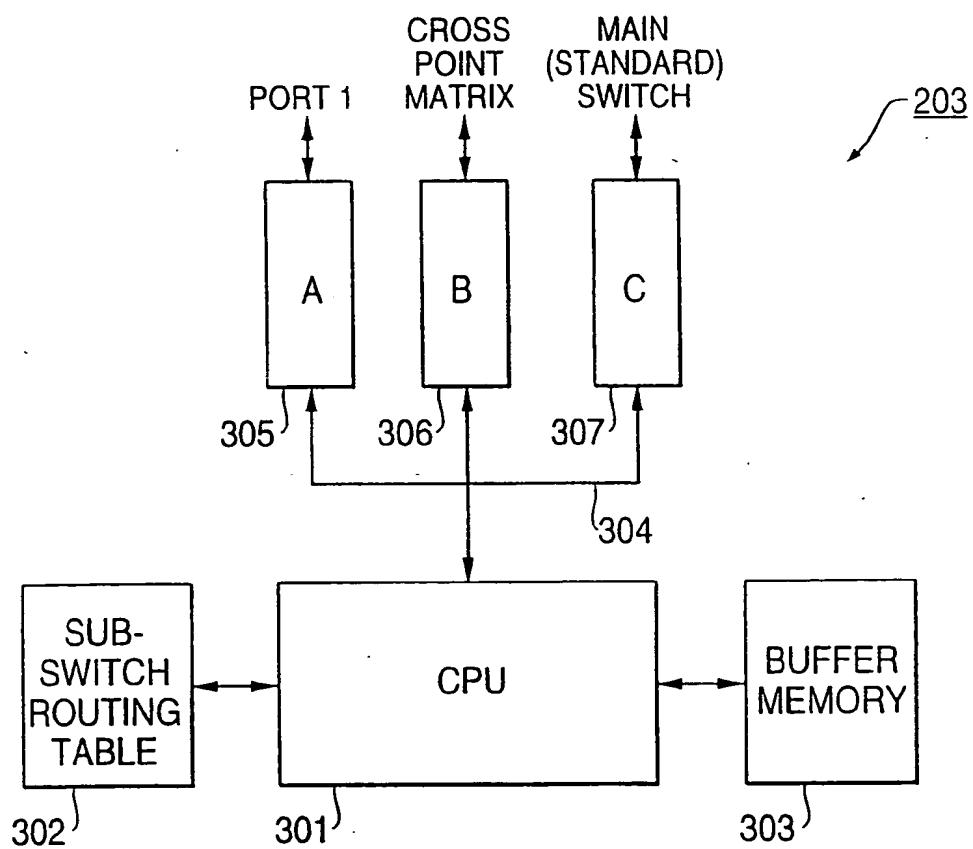
SUBSTITUTE SHEET (RULE 26)

FIG. 2



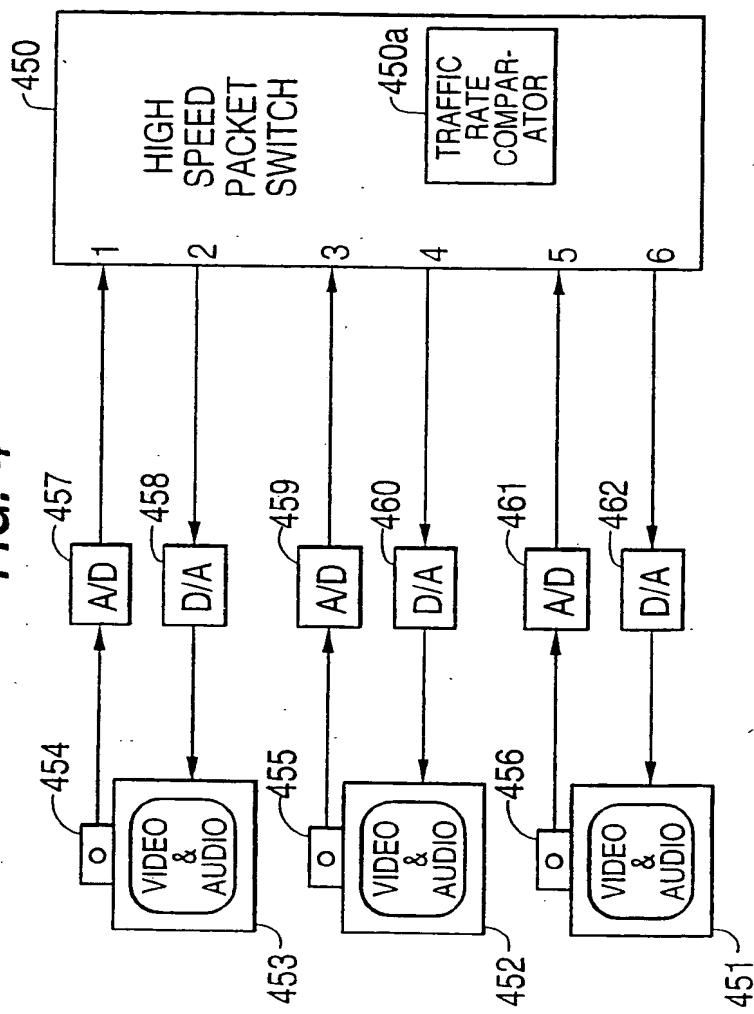
SUBSTITUTE SHEET (RULE 26)

FIG. 3



SUBSTITUTE SHEET (RULE 26)

FIG. 4



SUBSTITUTE SHEET (RULE 26)

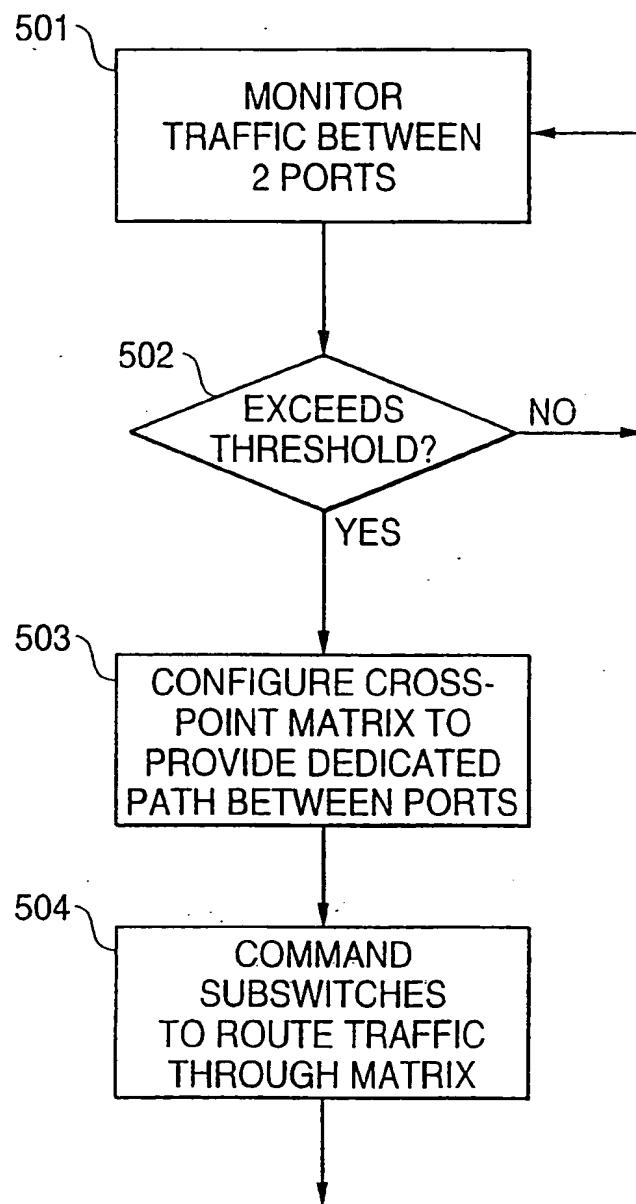
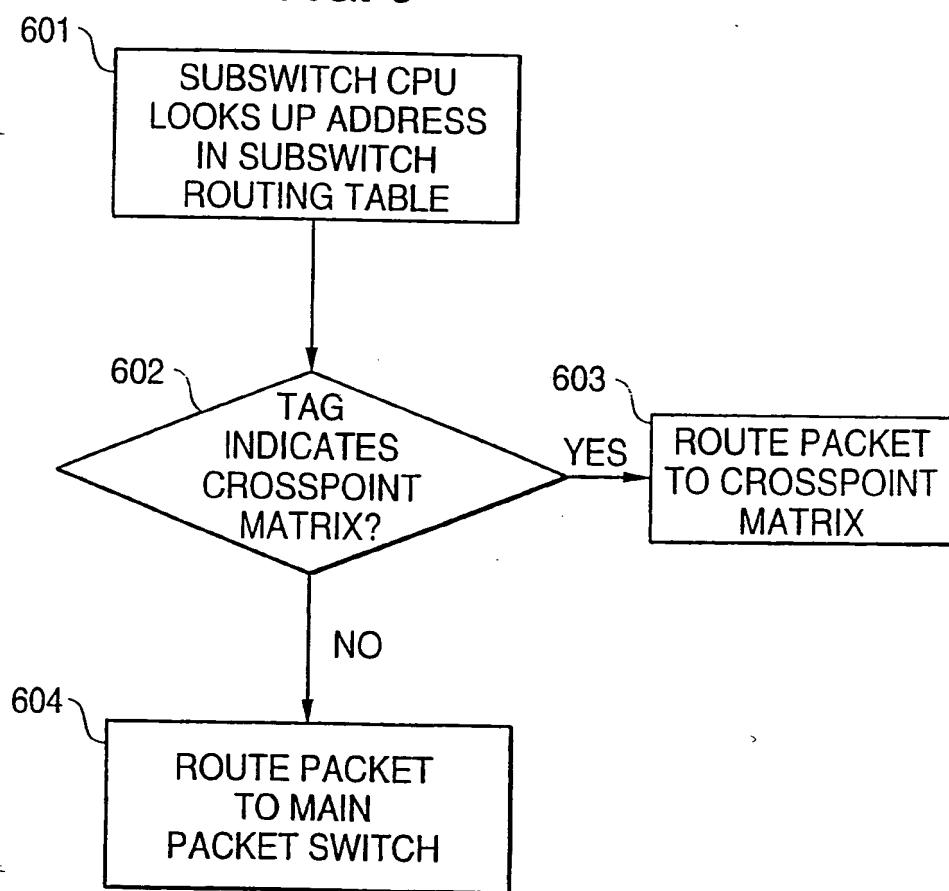
FIG. 5**SUBSTITUTE SHEET (RULE 26)**

FIG. 6



SUBSTITUTE SHEET (RULE 26)

INTERNATIONAL SEARCH REPORT

national Application No
PCT/US 8582A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04Q11/04 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04Q H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5 355 372 A (CHOWDHURY SHYAMAL ET AL) 11 October 1994 see abstract see figures 5,6,8 see column 1, line 64 - column 2, line 34 see column 3, line 16 - line 27 see column 5, line 28 - line 62 see column 6, line 38 - line 62 see claims 1,2 ---	1-5, 11-19
A	----- -/-	6-10,20

 Further documents are listed in the continuation of box C. Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the International filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "Z" document member of the same patent family

Date of the actual completion of the international search

29 January 1999

Date of mailing of the International search report

05/02/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Lamadie, S

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 98/18582

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	MANOJ KUMAR ET AL: "PERFORMANCE ENHANCEMENT IN BUFFERED DELTA NETWORKS USING CROSSBAR SWITCHES AND MULTIPLE LINKS" JOURNAL OF PARALLEL AND DISTRIBUTED COMPUTING, vol. 1, no. 1, 1 August 1984, pages 81-103, XP000084816 see figures 14,15 see page 98, paragraph 6 -----	1-5, 11-19
A	US 5 367 520 A (CORDELL ROBERT R) 22 November 1994 see abstract see figure 7 see column 4, line 4 - line 5 see column 4, line 27 - line 36 see column 5, line 60 - line 64 -----	6-10,20
A		1-20

INTERNATIONAL SEARCH REPORT

Information about patent family members

national Application No

PCT/US 91/0582

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
US 5355372	A 11-10-1994	JP	2546526 B	23-10-1996
		JP	7058754 A	03-03-1995
		JP	2546490 B	23-10-1996
		JP	6070350 A	11-03-1994
US 5367520	A 22-11-1994	NONE		

This Page Blank (uspto)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

This Page Blank (uspto)